

Mesh and hub-and-spoke networks on Azure

Architectural considerations

By Lamia Youseff and Nanette Ray
Azure Customer Advisory Team (AzureCAT)

December 2017

Contents

Introduction	3
Virtual network peering for LOB workloads	3
Mesh networks	5
Virtual network chaining: hub and spoke.....	6
Global mashup of mesh and hub-and-spoke network	7
Providing Azure services to business units.....	7
Architecture.....	7
Peering pros and cons	9
Controlling access through roles and policy.....	9
Learn more	10

List of figures

Figure 1. Typical LOB enterprise network configuration with Azure ExpressRoute connection to on-premises datacenter.	4
Figure 2. A mesh topology can be used to support a block chain consortium on Azure as long as the number of members is very small.	5
Figure 3. Peering connects one virtual network to another. It is not transitive—for example, LOB1 on the first hub cannot connect to LOB1 on the second hub.	6
Figure 4. An Azure customer's network topology for managed deployments.	8

Authored by Lamia Youseff and Nanette Ray (AzureCAT). Edited by RoAnn Corbisier.

© 2017 Microsoft Corporation. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Introduction

We are seeing a growing number of enterprises use [Azure Virtual Network Peering](#) to provide connected, secure workspaces for their business units. Virtual networks (the peers) are directly linked to each other using private IP addresses. The result is a low-latency, high bandwidth connection using the Microsoft backbone without the use of virtual private networks or gateways.

Virtual network peering gives Azure customers a way to provide managed access to Azure for multiple line-of-business (LOB) teams or to merge teams from different companies. But as enterprises expand their networks, they may encounter [subscription limits](#), such as the default maximum number of peering links. Most of these limits can be increased through a support request.

In this article, we look at the two main network topologies used by Azure customers.

Virtual network peering for LOB workloads

One Azure enterprise customer, typical of many we work with, supports multiple LOBs on Azure. Their IT group manages the process, provisioning Azure subscriptions for each business unit and enabling teams to work on their projects in relative isolation as well as keep track of costs. Each project commonly includes several environments such as development, production, and staging. This customer's security policy was to isolate the environments within their existing Azure Active Directory (Azure AD) tenant by providing each with its own virtual network as Figure 1 shows.

Their virtual network was composed of a front-end subnet, middle tier, and back-end subnet. Working with subnets gave the teams fine-grained control over their isolation boundaries and helped meet compliance requirements.

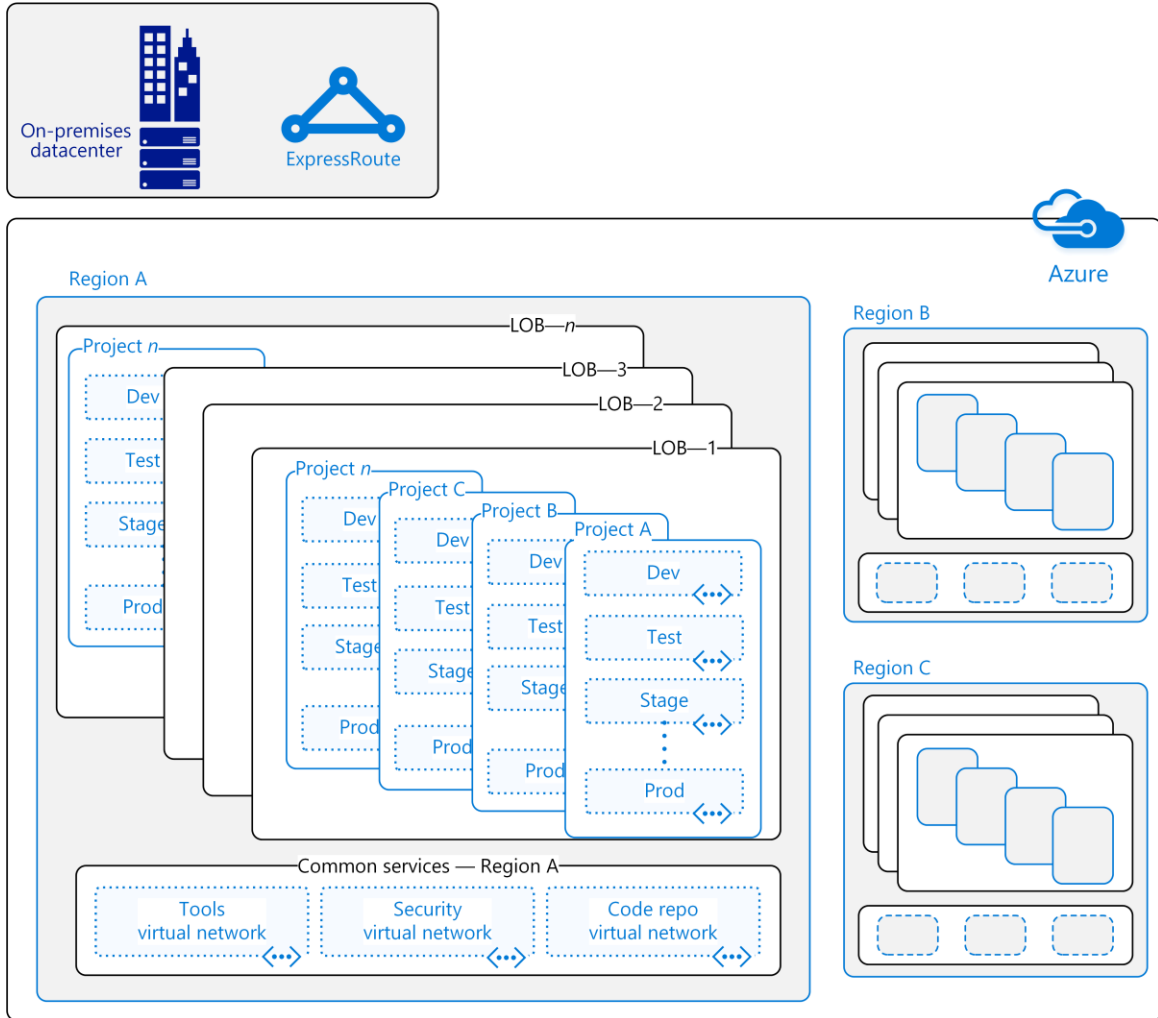


Figure 1. Typical LOB enterprise network configuration with Azure ExpressRoute connection to on-premises datacenter.

Very few projects run entirely in isolation. As Figure 1 shows, business units need access to:

- Other lines of business.
- A datacenter on premises, often through a shared ExpressRoute connection.
- A common services virtual network used to host services shared by the organization.

For example, in one Azure region, an enterprise might have dozens of LOBs, each with multiple virtual networks supporting development, testing, staging, and production environments. Yet they need to share source code repositories, developer tools, and security resources.

Using virtual network peering, these networks can be connected using either a mesh topology, ensuring all peers have access to all other peers, or a chaining topology to aggregate resources in hubs that can be shared by the spokes in the network.

Mesh networks

Most organizations solve their need for network isolation and connectivity by creating a mesh network architecture among the various virtual networks. All nodes in the network are interconnected, so network traffic is fast and can be easily redirected as needed.

The disadvantage of a mesh topology is that it requires many connections, making it costlier to operate and manage than other topologies. And the number of peering links required can quickly reach the Azure limits as Table 1 shows. The number of virtual network peering links is directly proportional to the number of virtual networks and business units.

Table 1. Proportional relationship of virtual networks to required peering links in a mesh topology.

Number of virtual networks	10	20	50	100
Peering links required	45	190	1,225	4,950

One Azure customer asked us for help creating a mesh network topology to support a block chain consortium. Block chains are distributed transactional systems that allow assets and investments to be exchanged based on smart contracts and predefined ledgers. In this scenario, four business partners wanted to start a *Société Anonyme* (SA), an anonymous corporation to serve as the basis of their block chain consortium. For block chain applications to run across the member’s digital assets, a mesh network topology provides connectivity between all virtual networks in one consortium as Figure 2 shows.

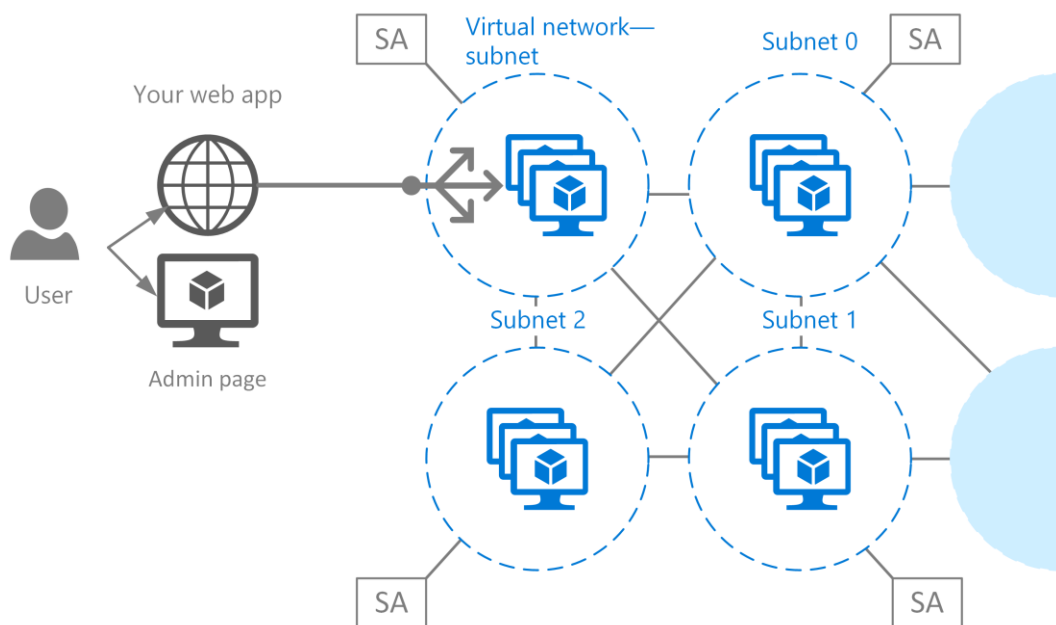



Figure 2. A mesh topology can be used to support a block chain consortium on Azure as long as the number of members is very small.

With four members, a mesh network requires three peering links per virtual network for a total of 12 links. With appropriate permissions, the members can peer virtual networks that exist in two different Azure AD tenants. The ability to set up peering links across Azure regions is in [preview](#), so the consortium could theoretically support global membership.

The consortium also plans to expand to at least 15 members, each with their own virtual network. When fully deployed, 14 peering links are needed per virtual network for a total of 105 links, but that might exceed the Azure defaults. So, they can then request an extension from Azure support. The real challenge arises when the planned expansion scales to 50 or more members. To support the new scale, virtual network chaining can provide an alternative architecture as the next section explains.

Virtual network chaining: hub and spoke

Another network topology commonly used by Azure customers relies on virtual network chaining, also known as hub and spoke. This topology connects groups of virtual networks to a hub virtual network that then connects to other hub virtual networks like links in a chain, as Figure 3 shows. When today's virtual network peering limitations make mesh networks problematic, a hub-and-spoke architecture offers an alternative that enables a larger number of virtual networks to be connected.



KEY:
 virtual network peering

Figure 3. Peering connects one virtual network to another. It is not transitive—for example, LOB1 on the first hub cannot connect to LOB1 on the second hub.

In this model, it's easy to add and remove spokes without affecting the rest of the network, giving business units great flexibility within their environments. This topology supports centralized monitoring and management. In addition, hubs aggregate connections so you're less likely to encounter the peering link limit. Azure also provides many ways to control network traffic using network security groups that specify rules to allow or deny traffic and user-defined routes that control the routing of packets.

Virtual network chaining supports a layered approach to isolation, which some large organizations require. For example, one multinational Azure customer wanted to provide greater security for their large-scale engineering projects, so they deployed one virtual network per project environment. A chaining model enables them to host each development, test, staging, and production environment in its own virtual network.

This topology does introduce a few tradeoffs. The hubs can become a potential single-point of failure—if one hub goes down for any reason, it breaks the chain. In addition, when communication between virtual networks travels through two or more hubs, the custom [route tables](#) and network virtual appliances (NVAs) can cause higher latency. The cost of these resources can add up, too, while the combination of traffic rules and peering links can make deployments challenging to deploy and manage.

For implementation details, see [Implement a hub-spoke network topology in Azure](#).

Global mashup of mesh and hub-and-spoke network

Another multinational professional services firm uses a hub and spoke architecture on Azure to support the work of projects deployed by multiple business units that access the same customer data. They conduct business in the United States, Europe, and Asia, where each region has its own datacenter. Projects may be deployed across Azure regions.

For example, a business unit might have a project to analyze global customer data. The team deploys their solution in their own virtual network spoke that is connected to the regional virtual network hub managed by the IT department. The hub network includes shared services—such as identity management—and is connected to the regional datacenter on premises, where the other centrally managed IT services reside.

One business unit had 300 active projects, and the company had to determine how to support virtual network peering without exceeding the limits.

Providing Azure services to business units

To make Azure services available to the business units, the company offers two provisioning models:

- **Managed:** The IT department provisions and manages the Azure subscription on behalf of the business unit. Like a traditional on-premises approach, a central IT team takes responsibility for deploying and maintaining the business unit's systems.
- **Unmanaged:** Self-sufficient teams can provision and own their Azure subscription. In this model, the business unit decides how to deploy, operate, and manage their solutions. Core components such as identity and network topology are shared.

Architecture

To support their provisioning model, the company created a global Azure topology, focusing on global network layout. Both managed and unmanaged deployments use the same topology. Each region has two hubs, one for managed deployments and the other for unmanaged. ExpressRoute connections link regions and are shared by the hubs and spokes. See Figure 4.

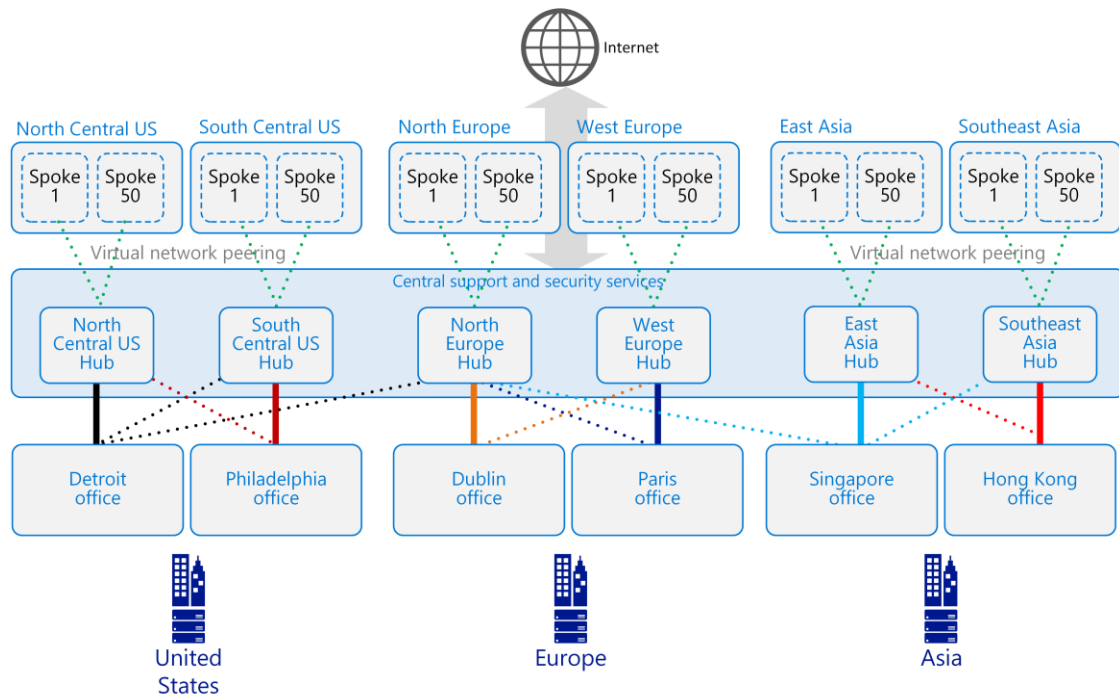


Figure 4. An Azure customer's network topology for managed deployments.

Components of this architecture include:

- Deployments across three Azure regional pairs (North Central US and South Central US, North Europe and South Europe, East Asia and Southeast Asia).
- Two dedicated ExpressRoute virtual circuits per continent, and two regional cross-connections per continent (North America, Europe, and Asia).
- Two virtual network hubs per Azure region (one per deployment type), each with up to 80 spokes connected by virtual network peering.

With this architecture, the company can constrain their virtual network peering links to the accepted levels. For example, the US North region has 15 managed virtual networks and 15 unmanaged. In Europe, only managed virtual networks are deployed, seven in the North Europe region and four in the West Europe region.

One thing to be aware of is that peering enables full communication by default. Sometimes our customers expect subnets to be unable talk to each other without explicit routes. When they implement a hub-and-spoke topology on Azure, they typically implement measures to control communications between network partners using [network security groups](#) and [user-defined routes](#).

Peering pros and cons

Peering provides many benefits with a few tradeoffs. On one hand, peered networks support isolation that frees business units to do their own thing while providing access to shared resources. It's easy to set up Azure virtual network peerings using Azure portal, PowerShell, or CLI, and there's no VPN gateway to configure.

On the other hand, growing networks can encounter peering limits pretty quickly. Azure supports a limited number of links by default, but you can request more from Azure support. Management can also grow more complex as organizations define a web of security policies and custom network settings such as [user-defined routes](#).

Other prerequisites apply:

- The peered virtual network cannot have overlapping IP address spaces. Inbound and outbound data transfer in the virtual network is charged at both ends of the peered networks.
- Peering connects one virtual network to another. It's not a pipeline where if A connects to B, and B to C, then A has a connection to C. Peering is one to one. A combination of user-defined routes and network security group configurations are needed to [build the chaining model](#).
- Classic Azure Service Manager deployments support peering with an Azure Resource Manager virtual network. Virtual networks within a classic Azure deployment can't be peered to each other. (Instead, use Azure [VPN Gateway](#).)
- In the block chain example, a question can arise about identity management if each member belongs to a different company, each with its own Azure architecture. Peering works only when the Azure subscriptions are associated with the same Azure AD tenant. If each member of the block chain consortium works for a different company, and each company is the tenant of a separate Azure subscription with Azure AD, the network would need to be set up through a [VPN Gateway](#) or some other way.

Controlling access through roles and policy

Some large Azure customers want to lock down isolation boundaries for greater security and control across an ecosystem of tens of thousands of virtual machines. In order to achieve that, they set up Azure Role-Based Access Control (RBAC) and define Azure Resource Manager policies at the project level.

RBAC defines the scope of user activities within Azure by segregating duties and granting access to resources based on need. For example, one role can be permitted to manage the virtual network and everything within it, while another role might manage only a single resource such as a SQL Database. Similarly, Resource Manager policies set limits on Azure resources and support deployment conventions. An operations team, for instance, might specify that only certain types of virtual machines or subnets are allowed.

Azure customers use a combination of RBAC and policies to reduce the number of virtual network peering links required in this topology so they don't exceed the limit. This approach also has the benefit of familiarity as seasoned security and operations teams have experience establishing policy and audit models.

Learn more

This article provides a high-level look at network topologies used by our customers. For implementation information, a good place to start is the Azure Reference Architectures website. For example, see:

[Implement a hub-spoke network topology in Azure.](#)

[DMZ between Azure and your on-premises datacenter](#)

[DMZ between Azure and the Internet](#)

The following articles also provide implementation details:

[Virtual network peering](#)

[Create a virtual network peering - Resource Manager, same subscription](#)

[Virtual network traffic routing](#)

[Filter network traffic with network security groups](#)

[Microsoft Azure Virtual Datacenter](#)

[Virtual Network pricing](#)